CENTRO STUDI IN ECONOMIA E FINANZA

CENTRE FOR STUDIES IN ECONOMICS AND FINANCE

WORKING PAPER no. 19

Consumption Insurance or Consumption Mobility

Tullio Jappelli and Luigi Pistaferri

April 1999

DIPARTIMENTO DI SCIENZE ECONOMICHE - UNIVERSITÀ DEGLI STUDI DI SALERNO
Via Ponte Don Melillo - 84084 FISCIANO (SA)
Tel. 089-96 3167/3168 - Fax 089-96 3169 – e-mail: csef@xcom.it

CSEF WORKING PAPER no. 19

Consumption Insurance or Consumption Mobility[*]

Tullio Jappelli[**] **and Luigi Pistaferri**[***]

Abstract

The theory of full consumption insurance posits that households are insulated from all idiosyncratic shocks so that the ratio of the marginal utilities of consumption of any two households is constant over time. Consumption insurance therefore implies absence of consumption mobility between any two time periods. This implication requires knowledge of the evolution of the entire consumption distribution, not just its mean as in standard tests of complete markets. We test this unexplored prediction of the theory using a panel drawn from the Bank of Italy Survey of Household Income and Wealth. We design an appropriate non-parametric test and find substantial mobility of consumption even controlling for possible preference shifts and measurement error in consumption. The findings strongly reject the theory of full consumption insurance.

Keywords: Consumption insurance, mobility

JEL classification: D52, D91, I30

[*] We thank Chris Carroll and seminar participants in Tilburg for helpful comments. This paper is part of a research project on "Structural Analysis of Household Savings and Wealth Positions over the Life Cycle". Financial support was provided by the Training and Mobility of Researchers Network Program (TMR) of the European Commission DGXII, the Italian National Research Council (CNR), the Italian Ministry for Universities and Scientific Research (MURST) and the James Lanner Memorial Fund

[**] CSEF, University of Salerno, and CEPR.

[***] University College London and CSEF, University of Salerno

Table of contents

1 Introduction

A large body of literature in industrialized and developing countries alike has proposed tests of full or partial consumption insurance (Cochrane, 1991; Townsend, 1994). The main implication of full consumption insurance is that the cross-sectional distribution of consumption over any group of households is constant over time. Therefore under complete markets, consumption growth is uncorrelated with changes in individual endowments. Of course aggregate consumption can increase or decrease, so that consumption growth for any household can be positive or negative, but the relative position of each individual in the cross-sectional distribution is preserved both in the short and the long run. Consumption insurance thus implies strong predictions about the entire consumption distribution, not just its mean or variance.[1]

In particular, the theory implies the total absence of consumption mobility between any two time periods, a much stronger proposition than is usually addressed by tests of consumption insurance. It follows that if one observes individuals moving up and down in the consumption distribution one must conclude that some people are not insulated from idiosyncratic shocks, which contradicts the assumptions of full consumption insurance. Although this implication of consumption insurance was mentioned in a theoretical paper by Banerjee and Newman (1991), to our knowledge it has never been explored in empirical analysis.

To test for the invariance of the consumption distribution one needs panel data. We construct a transition matrix for the distribution and apply non-parametric statistical tools to test the hypothesis of absence of consumption mobility between time periods. The empirical analysis is conducted on a panel of households drawn from the Bank of Italy Survey of Household Income and Wealth for the years 1987 to 1995.

With respect to previous studies that found overwhelming evidence against full consumption insurance (Cochrane, 1991; Attanasio and Davis, 1996) our contribution relates to both method and substance. On the methodological side, we analyze the transition matrix for household consumption and can therefore characterize the entire distribution of consumption rather than just its mean. Since we use a non-parametric index of market completeness, the statistical procedure is not sensitive to the particular utility function used, e.g. relative or absolute risk aversion.

[1] Deaton and Paxson (1994) show that the certainty equivalence version of the permanent income hypothesis implies that the cross-sectional dispersion in consumption of any given cohort should increase over time. They also note that full consumption insurance implies that the cross-sectional variance of consumption of the same cohorts should be constant over time.

On substance, examining the entire consumption distribution avoids arbitrary identifying assumptions. In fact, the statistical tests of consumption insurance used so far are tightly parametrized. To test the prediction that idiosyncratic shocks are uncorrelated with consumption growth, they rely on univariate regressions of consumption growth on aggregate variables and idiosyncratic shocks (such as change in household resources, unemployment hours, days of illness, etc.). Finding appropriate and exogenous proxies for the shocks is difficult in the extreme. Our procedure has several advantages: (i) we need not rely on any parametrized form for the utility function, (ii) we need not assume to identify any of these shocks; (iii) we need not assume that they are uncorrelated with unobservable or omitted preference shocks, including household fixed effects. Furthermore, the statistical test naturally provides an index of market completeness that measures the deviation of the actual consumption distribution from the distribution predicted by complete markets. This index can be used to compare the evolution of market completeness over time and check whether different population groups experience different degrees of consumption mobility.[2] Such information can be important for policy purposes. Consider for instance the possibility of a switch from a less to a more redistributive tax system and recall that tax progressivity provides implicit insurance to consumers. The effect of such policy change depends upon the amount of risk sharing already available in the economy. If private insurance markets are absent or largely incomplete, the policy change we are examining generates a welfare gain because it provides consumers with additional insurance; however, if consumers can fully insure the idiosyncratic shocks they face through private insurance markets, the policy change plays no role. Such a policy may even turn into a welfare loss if the provision of public insurance through progressive taxes crowds out private insurance schemes.

In Section 2 we review the model of consumption insurance and set out the basic intuition underlying our procedure. Section 3 presents the nonparametric test of consumption insurance and the mobility index. In Section 4 we explore the robustness of the test with respect to preference specification of the utility function and measurement error in consumption. The data and the empirical results are presented in Sections 5 and 6, respectively. We strongly reject full consumption insurance, in both the short and the long run and for each sample group that we analyze. The rejection of consumption insurance is not due to preference specification or measurement error. Section

[2]Hayashi, Altonji and Kotlikoff (1996) recommended that "future research should be directed to estimating the extent of consumption insurance over and above self-insurance" (p. 290). This paper is a step in this direction.

7 concludes.

2 Consumption insurance

First, let us review the model's main insights. Our argument does not rest on the specific form of the utility function; however, as a matter of convenience, we proceed on the assumption that households have identical preferences of the *CRRA* type, $u(c) = (1-\gamma)^{-1}c^{1-\gamma}$. If the social planner maximizes a weighted sum of individual households' utilities, the Lagrangian of the problem can be written (Deaton, 1997):

$$L = \sum_h \lambda_h \sum_s \sum_t \pi_{s,t} u(c_{h,s,t}) + \sum_s \sum_t \mu_{s,t} \left(C_{s,t} - \sum_h c_{h,s,t} \right)$$

where h, s and t are subscripts for the household h in the state of nature s in period t, λ_h is the social weight for household h, $\mu_{s,t}$ is the Lagrange multiplier associated with the resource constraint, $\pi_{s,t}$ the probability of state s in time period t, and $C_{s,t}$ aggregate consumption in state s and time t.

The first order condition can be written in logarithms as:

$$-\gamma \ln c_{h,s,t} = \ln \mu_{s,t} - \ln \lambda_h - \ln \pi_{s,t} \tag{1}$$

To obtain the rate of growth of consumption, one subtracts side-by-side from the expression at time $t+1$:

$$\Delta \ln c_{h,t+1} = -\gamma^{-1} \Delta \ln \mu_{t+1} + \gamma^{-1} \Delta \ln \pi_{t+1} \tag{2}$$

where we drop the subscript s because only one state is realized in each period. The two terms on the right-hand-side of equation (2) represent aggregate effects. The first is the growth rate of the Lagrange multiplier, the second is the growth rate of the state probabilities. Note that first-differencing has eliminated all household fixed effects.

Equation (2) states that the rate of growth of consumption of each household is the same. This implies that the initial cross-sectional distribution of consumption levels is a sufficient statistic to describe all future distributions: since all households have the same rate of growth of consumption, their relative position is stationary. Note that the stationarity of the cross-sectional distribution is directly implied by the assumption that insurance markets fully insulate households from idiosyncratic shocks.

The statistical counterpart of consumption insurance is that the transition matrix for household consumption is an identity matrix. In the next section we show how to construct such a transition matrix and how the matrix can be summarized by an appropriately designed mobility index. This index can be used to test the null hypothesis of no consumption mobility, which is implied by the theory of consumption insurance.

3 Consumption mobility

To summarize the transition matrix for consumption through an appropriate index of mobility, we build on an approach proposed by Shorrocks (1978). Assume that $\mathbf{P}$ is an unobservable $q \times q$ stochastic transition matrix of household consumption, q being the number of quantiles in the distribution. For notational simplicity we consider transition probabilities from period t to period $t+1$; extending the argument to transition probabilities in periods $t+2$, $t+3$, and so on, is straightforward. The generic element of $\mathbf{P}$ is p_{ij}, the probability of moving from quantile i in period t to quantile j in period $t+1$. Define n_{ij} as the number of households that move from quantile i in period t to quantile j in period $t+1$ and $n_i = \sum_i n_{ij}$ as the total number of observations in each row i of $\mathbf{P}$. The maximum likelihood estimator of the first-order Markov transition probabilities is $\widehat{p}_{ij} = \frac{n_{ij}}{n_i}$. The Shorrocks index of mobility is then defined as:[3]

$$S(\mathbf{P}) = \frac{q - trace(\mathbf{P})}{q} \tag{3}$$

If the probability of being in quantile i in period t is independent of that of being in quantile j in period $t+1$, the typical entry of the transition matrix is $p_{ij} = q^{-1}$ for all i and j. It follows that $trace(\mathbf{P}) = 1$ and $S(\mathbf{P}) = (q-1)/q$. Positing consumption insurance, the probability of being in quantile i in period t equals the corresponding probability in period $t+1$ and the probability of moving to a different quantile is zero. In this case the transition matrix is an identity matrix:

$$p_{ij} = \begin{cases} 0 & \text{if } i \neq j \\ 1 & \text{if } i = j \end{cases}$$

so that $trace(\mathbf{P}) = q$ and the index reaches its lower bound, $S(\mathbf{P}) = 0$. Since

[3] In its original formulation, the index is divided by $(q-1)$ rather than by q. We use this slight modification to bound the index between 0 and 1.

$0 \leq trace(\mathbf{P}) \leq q$, the mobility index satisfies the condition $0 \leq S(\mathbf{P}) \leq 1$.[4] $S(\mathbf{P})$ can be interpreted as the proportion of households moving across the consumption distribution between t and $t+1$.

The central limit theorem implies that $trace(\widehat{\mathbf{P}}) \overset{a}{\sim} N\left(\sum_i p_{ii}; \sum_i \frac{p_{ii}(1-p_{ii})}{n_i}\right)$, so that $S(\widehat{\mathbf{P}})$, the maximum likelihood estimator of $S(\mathbf{P})$, is asymptotically normally distributed (Schluter, 1998):

$$S\left(\widehat{\mathbf{P}}\right) \overset{a}{\sim} N\left(\frac{q - \sum_i p_{ii}}{q}; \ \frac{1}{q^2} \sum_i \frac{p_{ii}\left(1 - p_{ii}\right)}{n_i}\right)$$

Therefore one can test the null hypothesis of full consumption insurance, $S(\mathbf{P}) = 0$, using the statistic:

$$Z_1 = \frac{\frac{q - \sum_i \widehat{p}_{ii}}{q}}{\sqrt{\frac{1}{q^2} \sum_i \frac{\widehat{p}_{ii}\left(1 - \widehat{p}_{ii}\right)}{n_i}}} \sim N(0,1) \qquad (4)$$

The test is simple and powerful: the data requirements are minimal, because only the consumption distribution has to be known, and there is no need to identify exogenous idiosyncratic shocks. An important advantage is that the test does not rely on any specific form for the utility function. As the ordering of household consumption is invariant to monotonic transformation of the utility function, so are quantile probabilities.

It is often claimed that some population groups are more insulated than others from idiosyncratic shocks, or that households are more protected from such shocks in some periods than in others. To assess whether consumption mobility differs statistically over time or between population groups one can construct a test of differential mobility between two groups or time periods, based on the statistic:

$$Z_2 = \frac{S(\widehat{\mathbf{P}}_d) - S(\widehat{\mathbf{P}}_k)}{\sqrt{s.e.(S(\widehat{\mathbf{P}}_d))^2 + s.e.(S(\widehat{\mathbf{P}}_k))^2}} \sim N(0,1) \qquad (5)$$

where d and k are appropriately defined to allow comparisons over time or between population groups. Under the null hypothesis of no differential mobility, the statistic (5) is also asymptotically distributed as a standard normal.

[4] The upper bound is a case in which all households move to a different quantile so that $trace(\mathbf{P}) = 0$ and $S(\mathbf{P}) = 1$.

4 Extensions

The mobility index is derived assuming that the utility function is the same for all households and that there is no measurement error in consumption. In practice the index could potentially be upward biased by idiosyncratic preference shifts, preference heterogeneity or reporting errors. Supposing that demographic variables, household composition and labor supply affect marginal utility and not just consumption, the latter might rise or fall as these variables change over time. Part of the change in the consumption distribution as measured by the mobility index may therefore reflect genuine choices by households rather than uninsurable shocks. Likewise, consumption trajectories may differ because people have different preference parameters.

Measurement errors too can produce apparent consumption mobility. If households report their consumption with errors, one will find units moving up and down even with consumption insurance; hence, the index will tend to report higher mobility. We address these two problems in turn.

4.1 Preference specification and heterogeneity

Equation (1) suggests that the ratio between the marginal utilities of households h and h' is stationary. This does not always imply that the ratio of consumption levels is stationary, nor that the growth rate of consumption is the same for all households. Consider a case where the isoelastic utility function is augmented by a multiplicative preference shift θ:

$$u(c,\theta) = \theta \frac{c^{1-\gamma}}{1-\gamma}$$

It can be immediately shown that the growth rate of consumption for household h can be written as:

$$\Delta \ln c_{h,t+1} = -\gamma^{-1}\Delta \ln \mu_{t+1} + \gamma^{-1}\Delta \ln \pi_{t+1} + \gamma^{-1}\Delta \ln \theta_{h,t+1} \tag{6}$$

Equation (6) states that, over and above the effect of aggregate components, part of the cross-sectional movement in consumption growth is due to household-specific preference shifts (with the arrival of children, changes in household composition, age, and so on). If θ changes over time, the consumption distribution will no longer be stationary and the mobility index will be greater than zero even under consumption insurance. In the empirical section we therefore check for the robustness of the mobility index using *per capita* consumption and consumption *per adult equivalent*; we also experiment with a measure of consumption adjusted by a larger set of preference shifts.[5]

A related problem is the possibility that consumption and leisure may not be separable.[6] Although the implications of consumption insurance are unaffected when consumption and leisure are not separable, the right-hand-side of equation (2) includes another term, the rate of growth of the Lagrange multiplier of aggregate leisure. Cochrane (1991) points out that this term will vary between individuals except under the highly unrealistic assumption that the planner can freely transfer leisure across households. If the assumption is discarded, standard tests and our own procedure produce spurious evidence against consumption insurance. To address this problem, in the empirical section we augment the vector of preference shifts with the household head's leisure.

Note that our test is asymmetrically robust. The absence of consumption mobility (a result that does not reject consumption insurance) must imply that the preference shifts for which we do not control are not important determinants of the growth of marginal utility. In other words, the lack of consumption mobility cannot reflect estimator bias, as in more standard tests of the theory. Moreover, our test is robust in circumstances in which standard tests are not. The latter rely on regressions of consumption growth on idiosyncratic shocks and reject consumption insurance when the coefficients of the shock variables are significantly different from zero. But if the shocks are affected by measurement error, the OLS estimates are biased towards zero, providing spurious evidence in favor of consumption insurance. In contrast, our index will still report mobility because it does not require identifying idiosyncratic shocks in the first place.

An alternative way of introducing heterogeneity is to assume that the (unobservable) parameters of the utility function, say the degree of relative risk aversion, vary across individual:

$$u(c) = \frac{c^{1-\gamma_h}}{1-\gamma_h}$$

implying that the growth rate of consumption for household h is:

$$\Delta \ln c_{h,t+1} = -\gamma_h^{-1} \left(\Delta \ln \mu_{t+1} - \Delta \ln \pi_{t+1} \right) = -\gamma_h^{-1} \cdot \kappa_{t+1}$$

[5] In conventional tests of consumption insurance, preference shifts pose a rather different problem. If idiosyncratic shocks are correlated with omitted preference shifts, the estimated coefficients of the shock variables are biased, the direction of the bias depending on the correlation between preferences and shocks.

[6] We do not focus on non-separabilities between different consumption goods because in the empirical application we use a measure of total non-durable consumption.

Substituting in the expression above the growth rate of consumption in period t:

$$\Delta \ln c_{h,t+1} = \frac{\kappa_{t+1}}{\kappa_t} \cdot \Delta \ln c_{h,t}$$

Even if individual growth rates may be different, the period t ordering of growth rates will be identical in period $t+1$. While preference homogeneity implies that the initial cross-sectional distribution of consumption levels is a sufficient statistic for all future distributions, preference heterogeneity implies that the initial cross-sectional distribution of consumption growth rates is a sufficient statistic for all future distribution of growth rates implying that one should not observe mobility in the transition matrix for consumption growth between period t and $t+1$.

4.2 Measurement error

In the absence of preference shifts, consumption insurance delivers the following transition rule for *true* log-consumption:

$$\ln c_{h,t+1} = m_{t+1} + \ln c_{h,t} \tag{7}$$

where $m_{t+1} = -\gamma^{-1}\left(\Delta \ln \mu_{t+1} - \Delta \ln \pi_{t+1}\right)$. Now suppose that consumption is measured with a multiplicative error:

$$\ln c^*_{h,t+1} = \ln c_{h,t+1} + v_{h,t+1} \tag{8}$$

$$\ln c^*_{h,t} = \ln c_{h,t} + v_{h,t} \tag{9}$$

where $\ln c^*$ is *measured* consumption and v is a classical measurement error satisfying the assumption $v \sim i.i.n.d.\,(0, \sigma_v^2)$ (for simplicity, we also assume that its distribution is stationary). The transition law for log-consumption can be rewritten as:

$$\ln c^*_{h,t+1} = m_{t+1} + \ln c^*_{h,t} + v_{h,t+1} - v_{h,t}, \tag{10}$$

which implies that individual consumption growth is no longer a constant but an $MA(1)$ process with a time-varying drift, m_{t+1}. Measurement error therefore biases the mobility index $S(\widehat{\mathbf{P}})$ upwards: rejecting the null hypothesis $S(\mathbf{P}) = 0$ no longer implies that consumption insurance is violated.

The bias can be handled by noting that measurement error effectively increases the lower bound of the "true" mobility index $S(\mathbf{P})$. To see why, note first that regardless of consumption insurance the cross-sectional mean of $\ln c^*$ equals that of $\ln c$ because measurement errors average out. Note also that the difference between $var(\ln c^*)$ and $var(\ln c)$ depends on the variance of the measurement error. Since $\ln c^* = \ln c + v$, it follows that $var(\ln c^*) = var(\ln c) + \sigma_v^2$, or $\sigma_v^2 = \left[1 - \frac{var(\ln c)}{var(\ln c*)}\right] var(\ln c*) = \alpha \cdot var(\ln c^*)$.

The parameter α indicates the fraction of the cross-sectional variance of measured consumption that is contaminated by measurement error, ranging from 0 in absence of measurement error to 1 when the variance of measured consumption is entirely explained by measurement error. To get a feeling for how measurement error affects the statistical test, we use the variance-covariance matrix of consumption growth to estimate realistic values for α. We then perform a Montecarlo simulation under the null hypothesis of consumption insurance and measurement error. For each value of α we show how to generate different lower bounds of mobility, and then compare the actual mobility index with the theoretical index obtained under the joint hypothesis of consumption insurance and measurement error.

5 The data

The statistical test requires panel data on consumption. We use the 1987-1995 panel of the Italian Survey of Household Income and Wealth (SHIW). The dataset contains measures of consumption, income, and demographic characteristics of households. The SHIW provides a measure of total non-durable consumption, not just food, thus overcoming one of the main limitations of other panels, such as the PSID, that have been used to test for consumption insurance.

The SHIW is conducted by the Bank of Italy which surveys a representative sample of the Italian resident population. Sampling is in two stages, first municipalities and then households. Municipalities are divided into 51 strata defined by 17 regions and 3 classes of population size (more than 40,000, 20,000 to 40,000, less than 20,000). Households are randomly selected from registry office records. From 1987 through 1995 the survey was conducted every other year and covered about 8,000 households, defined as groups of individuals related by blood, marriage or adoption and sharing the same dwelling. Starting in 1989, each SHIW has reinterviewed some households from the previous surveys. The panel component has increased over time: 15 percent of the sample was reinterviewed in 1989, 27 percent in 1991, 43 percent in 1993, and 45 percent in 1995.[7] The net response rate (ratio of

[7] In the panel component, the sampling procedure is also determined in two stages: (i) selection of municipalities (among those sampled in the previous survey); (ii) selection of households reinterviewed. This implies that there is a fixed component in the panel (for

responses to contacted households net of ineligible units) was 64 percent in 1987, 38 percent in 1989, 33 percent in 1991, 58 percent in 1993, and 57 percent in 1995. Details on sampling, response rates, processing of results and comparison of survey data with macroeconomic data are provided by Brandolini and Cannari (1994).[8]

To minimize measurement error we exclude cases in which the head changes over the sample period or gives inconsistent age figures. The total number of transitions is 10,508. After the exclusions, the sample has 9,214 transitions. Table 1 reports sample statistics of log consumption and other household characteristics. All statistics are computed using sample weights. The panel is relatively stable over the sample period. Consumption grows considerably between 1987 and 1989 and is stable afterwards. Over time, family size declines while the number of income recipients increases. Other demographic characteristics remain roughly unchanged. The fall in self-employment is paralleled by an increase in public employees.

6 Empirical results

We first present full-sample results. We then address the issue of preference specification and measurement errors in consumption. Finally, we focus on consumption mobility in specific population groups.

6.1 Full sample estimates

There are two methods for constructing a transition matrix. One is to keep the width of the interval in which consumption is discretized constant and let the number of observations within each interval vary. The alternative is to keep constant the marginal probabilities and let the interval width change, for instance dividing the distribution into discrete quantiles. The second method is more standard. We proceed using quartiles throughout; results with deciles are qualitatively similar and are not reported. In what follows, we focus on the distribution of the logarithm of non-durable consumption, but results for consumption levels or for any monotonic transformation of consumption are the same.

instance, households interviewed 5 times between 1987 to 1995, or 4 times from 1991 to 1995) and a new component every survey (for instance, households reinterviewed only in 1989).

[8] In the panel section, the net response rate was 25 percent in 1989, 54 percent in 1991, 71 percent in 1993, and 78 percent in 1995. The lower attrition rates in 1991-1995 reflect the fact that participation was made voluntary after 1989. According to Bank of Italy statisticians the amount of attrition is relatively modest (Brandolini, 1998).

Table 2 reports the transition matrix pooling all transitions over all years. Recall that the SHIW is run every two years, so we observe transitions from period t to period $t+2$. The elements of the main diagonal report the proportion of households that did not change quartile. For instance, the entry in the top left of the table indicates that 66 percent of the households in the first quartile at time t were still in that quartile two years later. Off-diagonal elements signal consumption mobility. For instance, the second entry in the first row indicates that 25 percent of households moved from the first quartile in period $t-2$ to the second quartile in period t. Overall, the table shows quite a substantial amount of consumption mobility. About one third of households in the first quartile move upwards in the consumption distribution, about one third in the fourth quartile move downwards, and more than half in the third and fourth quartiles move either up or down.[9]

Further insights about the evolution of the cross-sectional distribution can be gained by examining the probability of households changing quartile in the sample period. In Figure 1 we denote these values as "mobility probabilities". The probability is relatively high for the second and third quartiles (about 60 percent) and lower in the top and bottom quartiles (between 30 and 40 percent). The figure indicates not only that there is substantial consumption mobility in all quartiles, but also that the mobility is persistent in all survey years. As we shall see, the results of the descriptive evidence are confirmed by the statistical test.

The mobility index corresponding to the elements of the matrix in Table 2 is reported in the first row of Table 3. The statistic has a value of 0.47, with a standard error of 0.005. The null hypothesis of consumption insurance, $S(\mathbf{P})=0$, is therefore overwhelmingly rejected. This is in line with previous studies for the United States that also reject the hypothesis. The other rows of Table 3 report mobility for selected periods, as the sample was marked by economic expansion in the early years and by the deep recession in 1991-93. Overall, no great variability in consumption mobility is indicated (the index ranges from 0.44 to 0.51). In the long run mobility is still as high as 0.40.[10]

The descriptive and statistical analyses suggest that between 1987 and

[9] The symmetry of the transition matrix can be tested using the maximum likelihood test suggested by Bishop, Fienberg and Holland (1988). The statistic is of the form $\Psi=\sum_{i>j}\frac{(p_{ij}-p_{ji})^2}{p_{ij}+p_{ij}}\sim\chi^2_{q(q-1)/2}$. The p-value of the test is close to 1, and does not reject the hypothesis that the transition matrix is symmetric.

[10] A Markov process is a stochastic process in which the probability of entering a certain state depends only on the previous state and on the matrix governing the process. If these assumptions hold for the stochastic transition matrix $\mathbf{P}$, it is possible to determine the limit (or long-run) state as the eigenvector of the matrix $\mathbf{P}$ associated with the eigenvalue 1.

1995 the Italian economy was characterized by substantial consumption mobility. On average, half of the population moved up or down in the distribution every two years, a result that strongly controverts full consumption insurance. The counterpart of this finding is that half of the households are unable to insure idiosyncratic shocks by formal or informal market arrangements. If consumption is regarded as a proxy for permanent income, it turns out that the latter is not so permanent after all.

As Deaton and Paxson (1994) note, consumption insurance implies that the cross-sectional variance of log-consumption will be constant over time. In our sample this hypothesis is not rejected (the p-value associated with this hypothesis is 0.73).[11] However, the stationarity of the cross-sectional variance does not imply absence of consumption mobility and cannot be cited as evidence for consumption insurance. Since the results indicate that consumption is mobile but that the variance is roughly constant, it must be that the variance of the cross-sectional distribution is not adequate to measure consumption mobility. This is one case in which simple measures of dispersion must be supplemented by careful analysis of the entire distribution.

6.2 Preference specification and heterogeneity

As we mention in Section 4.1, the marginal utility of consumption is likely to be affected by demographic or labor supply variables that change over time. In this case we would observe mobility even in the absence of non-insurable shocks. One of the most important demographic variables that can affect preferences is certainly household composition. For instance, the arrival of children alters family needs, hence consumption allocation. We thus compute mobility defining transitions in terms of *per capita* consumption and consumption per adult equivalent; the latter is more appropriate in the presence of economies of scale.[12] Table 4 shows, however, that using *per capita* consumption makes no difference with respect to Table 3 while the adult equivalent measure increases the mobility index only slightly.

Defining consumption per adult equivalent eliminates just one of the possible sources of predictable consumption variability. To take account of a larger set of demographic variables potentially affecting marginal utility, we can rewrite equation (6) as:

$$\Delta \ln c_{h,t} - \gamma^{-1} \Delta \ln \theta_{h,t} = m_{t+1} \tag{11}$$

[11] More precisely, this is the p-value of a test that $s.d.(\ln c_t) = s.d.(\ln c_{t-1})$.

[12] Adult equivalency is defined as: $1 + 0.8(\textit{Number of adults} - 1) + 0.25(\textit{Number of children})$. Data for 1987 are not used because information on the number of children is lacking.

where as before $m_{t+1} = -\gamma^{-1}\left(\Delta \ln \mu_{t+1} - \Delta \ln \pi_{t+1}\right)$. Equation (11) implies that the ratio of marginal utilities for any two households in the cross-section is stationary after controlling for preference shifts. Our procedure consists in two steps. First we impute a measure of consumption adjusted for demographic effects, $\ln \widetilde{c}_{h,t} = \ln c_{h,t} - \widetilde{\gamma}^{-1} \ln \theta_{h,t}$, where $\widetilde{\gamma}$ is the OLS estimate of a regression of $\ln c_{h,t}$ on $\ln \theta_{h,t}$. The θ variables that we use are family size, age, age squared, number of children and number of income recipients. This yields a measure of consumption in which demographic effects have been filtered out. In the second stage we construct transitions on the generated variable $\ln \widetilde{c}_{h,t}$ and test the absence of consumption mobility. The resulting index is again quite close to that estimated without controlling for demographic effects (0.53 with a standard error of 0.005).[13]

Leisure is another factor that might affect the marginal utility of consumption. In Figure 2 we plot the empirical distribution of annual working hours of household heads.[14] We find the expected concentration of observations at 0 (unemployment or retirement) and 2080 (a standard work week of 40 hours). The low variability of the distribution reflects the well-known rarity of part-time jobs in the Italian labor market. The limited flexibility of hours is *prima facie* evidence that changes in leisure should not be a major factor in explaining consumption mobility.

A more formal test of the effect that leisure has on consumption mobility consists in including leisure in the first-stage regression described above. In this case the mobility index increases to 0.60, whether log-leisure is instrumented with past values or not. If leisure were responsible for some of the consumption transitions one should observe a decline, not an increase, in the mobility index.[15] Therefore, we conclude from this section that mispecification of preferences explains little or none of the consumption mobility of our sample.

As a final check of the potential impact of preference heterogeneity on mobility, we construct a transition matrix for consumption growth rates.[16]

[13]We also experiment with a wider set of demographic variables (number of children in various age bands, education, region of residence, city size). The mobility index is virtually unchanged.

[14]The density function is estimated non-parametrically by a standard kernel method. We use the optimal bandwith suggested by Silverman (1986). The 1987 distribution is omitted because data on labor supply are not available .

[15]The mobility index can be biased downward if leisure or preference shifts increase consumption needs and if they are negatively correlated with the rank in the consumption distribution, i.e.if they affect more strongly households in the bottom part of the consumption distribution.

[16]This requires at least three years of observations. The sample size for this experiment is therefore reduced to 3,341 transitions.

The associated index is 0.81 with a standard error of 0.07, confirming that our sample displays substantial consumption growth mobility. The finding of higher mobility in growth rates than in levels suggests that also preference heterogeneity is unlikely to explain the rejection of consumption insurance. Therefore, we conclude from this section that mispecification or heterogeneity of preference explains little or nothing of the consumption mobility that we observe in the data.

6.3 Measurement error

In Section 4.2 we define α as the proportion of the variance of measured log-consumption due to measurement error. Clearly the bias in mobility increases with α. Here we provide evidence on the size of α and the likely impact of measurement error on the estimate of mobility; we also provide bounds of the estimator of mobility in the presence of measurement error in consumption.

Even in the presence of measurement error, complete markets impose strong restrictions on the covariance matrix of consumption. In fact, writing equation (10) as: $\Delta \ln c^*_{h,t+1} = m_{t+1} + \Delta v_{h,t+1}$, omitting the aggregate component, the following testable restrictions are implied by the autocovariance matrix of $(\Delta \ln c^*)$:

$$
\begin{aligned}
E\left[\left(\Delta \ln c^*_{h,\tau}\right)^2\right] &= 2\sigma_v^2 \\
E\left[\left(\Delta \ln c^*_{h,\tau}\right)\left(\Delta \ln c^*_{h,\tau-1}\right)\right] &= -\sigma_v^2 \\
E\left[\left(\Delta \ln c^*_{h,\tau}\right)\left(\Delta \ln c^*_{h,\tau-j}\right)\right] &= 0 \text{ for all } j \geq 2
\end{aligned}
$$

We are interested in identifying $\alpha = \frac{\sigma_v^2}{var(\ln c^*)}$. To estimate σ_v^2, we first define a mean zero measure of *per capita* log consumption adjusted for aggregate shocks:

$$\xi_{h,t} = \ln c^*_{h,t} - \ln \theta_{h,t} - \widetilde{\Gamma}_t$$

where θ includes only family size and $\widetilde{\Gamma}$ is the cross-sectional mean of consumption per capita. The covariance matrix of $\Delta\xi_{h,t}$ is given in Table 5. At first sight, the pattern is not inconsistent with the restrictions implied by consumption insurance and measurement error: the first order autocovariances are negative and statistically significant, second and higher order autocovariances are small, not statistically significant different from zero. Note also that the empirical pattern of autocovariances in Table 5 is inconsistent with persistent measurement error.

At face value, the covariance matrix suggests that the variance of measurement error is on the order of 0.06 (average over all years). Since the overall variance of consumption is about 0.29 (average over all years), measurement error explains roughly one fifth of the overall variance (α=0.06/0.29$\simeq$0.2). By comparison with the PSID, where researchers have found much larger estimates of α (between 70 and 90 percent), our covariance matrix suggests that the SHIW data on total non-durable consumption are of much better quality than the PSID data on food consumption.

Even though α=0.2 is not a high number, it must be regarded as an unlikely upper bound for the fraction of the variance explained by measurement error. Recall that this value is obtained on the hypothesis of full consumption insurance. Suppose, however, that consumption insurance does not hold and that an idiosyncratic shock $\eta_{h,t}$ affects consumption growth, $\Delta \ln c^*_{h,t+1} = m_{t+1} + \Delta v_{h,t+1} + \eta_{h,t}$. Assume that $\eta_{h,t}$ is uncorrelated with measurement error at all leads and lags. The restrictions on the covariance matrix of the adjusted measure of consumption growth can then be rewritten:

$$\begin{aligned} E\left[\left(\Delta \ln c^*_{h,\tau}\right)^2\right] &= 2\sigma_v^2 + \sigma_\eta^2 \\ E\left[\left(\Delta \ln c^*_{h,\tau}\right)\left(\Delta \ln c^*_{h,\tau-1}\right)\right] &= -\sigma_v^2 \\ E\left[\left(\Delta \ln c^*_{h,\tau}\right)\left(\Delta \ln c^*_{h,\tau-j}\right)\right] &= 0 \text{ for all } j \geq 2 \end{aligned}$$

Note that the restrictions now imply that the variance exceeds, in absolute value, twice the covariance, $E\left[\left(\Delta \ln c^*_{h,\tau}\right)^2\right] > -2E\left[\left(\Delta \ln c^*_{h,\tau}\right)\left(\Delta \ln c^*_{h,\tau-1}\right)\right]$.[17] A test that $E\left[\left(\Delta \ln c^*_{h,\tau}\right)^2\right] = -2E\left[\left(\Delta \ln c^*_{h,\tau}\right)\left(\Delta \ln c^*_{h,\tau-1}\right)\right]$ against the one-sided alternative $E\left[\left(\Delta \ln c^*_{h,\tau}\right)^2\right] > -2E\left[\left(\Delta \ln c^*_{h,\tau}\right)\left(\Delta \ln c^*_{h,\tau-1}\right)\right]$ rejects the null (the t-statistic is 2.75 with a p-value of 0.003).[18] The rejection is also apparent from the pattern of covariances reported in Table 5, particularly for 1989-91 and 1991-93. This example indicates that the autocovariance matrix is affected by something other than measurement error alone. To estimate the consumption variability that cannot be attributed to measurement error, note that:

$$\sigma_\eta^2 = E\left[\left(\Delta \ln c^*_{h,\tau}\right)^2\right] + 2E\left[\left(\Delta \ln c^*_{h,\tau}\right)\left(\Delta \ln c^*_{h,\tau-1}\right)\right] \tag{12}$$

[17]Equivalently, under the hypothesis maintained, this implies $\sigma_\eta^2 > 0$.

[18]This test is pooled over all years. For single years, the null hypothesis is rejected for 1991 and 1993 but not for 1995.

One possibility is to choose values of σ_η^2 that are consistent with a significance level of 5 percent or higher. In our sample the null hypothesis that $\sigma_\eta^2 = 0.025$ has a p-value of 0.049 and therefore cannot be rejected. This implies $\sigma_v^2 = 0.035$ (averaged over all years) and $\alpha = 0.12$. In more realistic examples, first-order autocovariances alone are not sufficient to identify σ_v^2, so that α is likely to be lower than 0.12. For instance, if the idiosyncratic shock η is persistent one cannot disentangle the fraction of the variance due to measurement error from that due to shocks.[19] From the foregoing, we conclude that α=0.12 is an upper bound to measurement error and that more realistic values of α range from 0.05 to 0.10.

The next step is to assess how measurement error affects the mobility index under the null hypothesis of consumption insurance. For this purpose, we design a Montecarlo simulation based on 100 replications, using *per capita* consumption throughout. In each year we choose a sample size identical to the number of transitions (for instance, it is 3,211 for 1993-95). Measurement errors at times t and $t-2$ are drawn from a normal distribution with mean zero and variance α times the variance of measured consumption at t and $t-2$. True consumption $\ln c_{t-2}$ is drawn from a normal distribution with mean equal to the mean of measured consumption and variance of $(1-\alpha)$ times the variance of measured consumption at $t-2$. Under the null hypothesis of consumption insurance, $\ln c_t = m_t + \ln c_{t-2}$, where m_t is the aggregate consumption growth, estimated as the average of individual consumption growth rates between $t-2$ and t. Given our assessment of the likely magnitude of measurement error, we choose values for $\alpha = \{0.05, 0.1, 0.12\}$ and simulate the mobility index $S(\widetilde{\mathbf{P}})$.The results of the simulation are reported in Table 6. The first column reproduces the actual mobility index $S(\widehat{\mathbf{P}})$ of consumption *per capita*, from Table 4, column 2. If α=0.05 the simulated index $S(\widetilde{\mathbf{P}})$=0.26 in 1987-89, against $S(\widehat{\mathbf{P}})$=0.47. The fraction of mobility that cannot be attributed to measurement error is $\frac{S(\widehat{\mathbf{P}})-S(\widetilde{\mathbf{P}})}{1-S(\widetilde{\mathbf{P}})}$=0.29. If α=0.10 this fraction is 0.19; even in the most unfavorable case of α=0.12 the fraction of "true" mobility is 0.15. To summarize, in 1987-89 the fraction of households that move across the consumption distribution for reasons other

[19] To take one example, if $\eta_{h,t} = \psi_{h,t} - \rho\psi_{h,t-1}$, the restrictions can be rewritten as:

$$\begin{aligned}
E\left[\left(\Delta \ln c_{h,\tau}^*\right)^2\right] &= 2\sigma_v^2 + (1+\rho^2)\sigma_\psi^2 \\
E\left[\left(\Delta \ln c_{h,\tau}^*\right)\left(\Delta \ln c_{h,\tau-1}^*\right)\right] &= -\sigma_v^2 - \rho\sigma_\psi^2 \\
E\left[\left(\Delta \ln c_{h,\tau}^*\right)\left(\Delta \ln c_{h,\tau-j}^*\right)\right] &= 0 \text{ for all } j \geq 2
\end{aligned}$$

and identification would no longer be possible.

than measurement error ranges from 15 at $\alpha = 0.12$ to 47 percent at $\alpha = 0$. Similar results are obtained for transitions in other years.[20]

6.4 Sub-sample estimates

Our statistical test allows us to inquire into population groups which are most exposed to idiosyncratic shocks. To check whether there are differences in consumption mobility we use the statistic on difference of means discussed in Section 3. Again, we use a measure of *per capita* consumption throughout.

Table 7 reports consumption mobility for households with different regions of residence, occupations (public vs. private and self-employed vs. employee), education, year of birth and number of income recipients. Mobility is greater in the North than in the South (0.50 against 0.48), a difference possibly explained by the greater social insurance role offered by the family and the presence of informal market arrangements in the South. Mobility is also higher in the private sector than in the public sector (0.49 against 0.45), a reflection of the fact that in Italy public sector employees enjoy stable earnings tied to strict seniority rules rather than performance, virtually job security. Employees, which face less income risk than the self-employed, also exhibit lower consumption mobility (0.47 against 0.50). The difference between households with only compulsory education and those with college degrees is not statistically significant. The comparison by year of birth suggests that younger cohorts are progressively more able to smooth away idiosyncratic shocks.

Common sense suggests that households with multiple earners can insure income shocks better than single earners. We distinguish three groups: those with no change in number of earners, with a decrease an with an increase. The results indicate that mobility is greatest among the latter two groups (0.55 and 0.56 respectively).

Overall, we find plausible and significant variation of mobility by occupation and demographic group. Some of the differences can be tied to specific hypothesis concerning the working of credit, insurance and informal markets. This is certainly the case for the relative low mobility of public sector employees and households where the number of income recipients does not change. However, contrary to our expectations, overall we find a surprisingly small amount of variability between different groups. Table 7 indicates that even if in most cases the mobility indexes are significantly different from each

[20] Seeking to minimize the impact of measurement errors by focusing on households with low rates of consumption growth ($-1 \leq \Delta \ln c \leq 1$) has virtually no effect on the results (the mobility index is 0.48).

other statistically, the p-values associated with the difference in means generally indicate marginal rejection of the null hypothesis of equality between groups. Furthermore, the differences in the mobility index between groups are generally not large in absolute value.

7 Conclusions

Consumption insurance implies that in any time period the initial cross-sectional distribution of consumption is a sufficient statistic for all future distributions. This implication of consumption insurance is as yet unexplored. We construct a transition matrix for total non-durable consumption using the 1987-95 panel contained in the Bank of Italy Survey of Household Income and Wealth. We then summarize the transition matrix of consumption by an appropriate mobility index. The test of consumption insurance we propose is simple and powerful. Most importantly, the non-parametric test proposed does not depend on functional form, identification assumptions about the source of idiosyncratic shocks, or their potential correlation with omitted preference shifts.

We find that roughly 50 percent of households move up or down in the consumption distribution between any two periods, in both the short and the long run. This constitutes very strong evidence against consumption insurance. There are interesting variations in the mobility patterns within different population groups, but overall the inter-group variation in mobility is not large. The mobility observed is unlikely to be explained by the effect of preference shifts. Consumption *per capita* and per adult equivalent exhibit mobility comparable to that of total non-durable consumption. When we control for other potentially important observable preference shifts (such as family size, age, education, and leisure) mobility actually increases. Finally, we find substantial mobility in consumption growth, not only in consumption levels, implying that preference heterogeneity does not explain rejection of consumption insurance.

Part of the consumption mobility observed in the sample may be due to measurement error. We show that in our data measurement error is unlikely to explain a large fraction of the total cross-sectional variance of consumption. To assess the impact of measurement error on the mobility index we then perform a Montecarlo experiment. The simulation shows that our test rejects the hypothesis of consumption insurance even in the most unfavorable case, one in which measurement error has the highest impact on mobility. We conclude that in Italy a great deal of consumption mobility is explained by idiosyncratic shocks that households are unable to insure. Even though our test is powerful disproof of consumption insurance, it does not constitute evidence either for or against the permanent income hypothesis; these "are distinct propositions, and each may hold independently of the other" (Cochrane, 1991). In future research we plan to use transition matrices to model consumption and income mobility jointly and look into their implications for smoothing idiosyncratic shocks across different states of nature.

References

[1] Attanasio, Orazio P., and Steve Davis (1996), "Relative wage movements and the distribution of consumption", *Journal of Political Economy,* **104**, 1227-62.

[2] Banerjee, Abhijit V. and Andrew F. Newman (1991), "Risk-bearing and the theory of income distribution", *Review of Economic Studies*, **58**, 211-35.

[3] Bishop, Y.M.M., S.E. Fienberg, and P.W. Holland (1988), *Discrete multivariate analysis,* MIT Press.

[4] Brandolini, Andrea (1998), "The personal distribution of income in post-war Italy: source description, data quality, and the time pattern of income inequality", mimeo, Bank of Italy.

[5] Brandolini, Andrea, and Luigi Cannari (1994), "Methodological Appendix: The Bank of Italy Survey of Household Income and Wealth", in *Saving and the accumulation of wealth. Essays on Italian households and government behavior*, edited by Albert Ando, Luigi Guiso and Ignazio Visco. Cambridge: Cambridge University Press.

[6] Cochrane, John (1991), "A simple test of consumption insurance", *Journal of Political Economy* **99**, 957-976.

[7] Deaton, Angus (1997), *The analysis of household surveys.* Baltimore: The Johns Hopkins University Press.

[8] Deaton, Angus, and Christina Paxson (1994), "Intertemporal choice and inequality", *Journal of Political Economy*, **102**, 384-94.

[9] Hayashi, Fumio, Joseph Altonji, and Laurence Kotlikoff (1996) "Risk sharing between and within families", *Econometrica* **64**, 261-94.

[10] Schluter, Christian (1998), "Statistical inference for mobility indexes", *Economic Letters*, **59**, 157-62.

[11] Shorrocks, Anthony F. (1978), "The measurement of mobility", *Econometrica*, **46**, 1013-24.

[12] Silverman, B.W. (1986), *Density estimation for statistics and data analysis,* Chapman and Hall.

[13] Townsend, Robert (1994), "Risk and insurance in village India", *Econometrica*, **62**, 539-91.

Table 1
Descriptive statistics

	1987	1989	1991	1993	1995	All years
$\ln c_t$	9.90	10.08	10.02	10.01	10.00	10.02
$var\,(\ln c_t)$	0.26	0.26	0.29	0.29	0.27	0.28
South	0.41	0.37	0.34	0.36	0.39	0.37
North	0.43	0.46	0.48	0.47	0.43	0.46
Family size	3.15	3.12	3.04	3.07	3.01	3.07
Self-employed	0.20	0.17	0.17	0.16	0.15	0.16
Public employee	0.17	0.18	0.27	0.23	0.25	0.23
Years of schooling	7.38	7.97	8.19	8.03	8.10	8.03
Born $\leq$ 1927	0.33	0.29	0.26	0.24	0.24	0.26
Born 1928-1937	0.20	0.21	0.20	0.18	0.19	0.19
Born 1938-1947	0.26	0.25	0.24	0.23	0.22	0.55
Income recipients	1.63	1.72	1.72	1.74	1.78	1.73
# of obs.	1,097	2,717	4,036	4,006	3,211	15,067

Note: Cross-sectional means and variances are computed using sample weights.

Table 2
The transition matrix of consumption

	Quartile at time t			
Quartile at time $t-2$	1^{st}	2^{nd}	3^{rd}	4^{th}
1^{st}	0.66	0.25	0.07	0.02
2^{nd}	0.24	0.41	0.26	0.09
3^{rd}	0.09	0.26	0.41	0.24
4^{th}	0.02	0.10	0.25	0.63

Note: The table reports consumption transitions from period t to period $t+2$. Transitions are pooled over all sample years.

Table 3
Mobility index

Panel	Number of transitions	$S\left(\widehat{\mathbf{P}}\right)$	$s.e.\left(S\left(\widehat{\mathbf{P}}\right)\right)$
All years	9,204	0.4729	0.0051
1987-1989	1,097	0.5066	0.0146
1989-1991	1,914	0.4621	0.0110
1991-1993	2,982	0.5060	0.0090
1993-1995	3,211	0.4367	0.0085

Note: The table reports the mobility index computed using the transition matrix pooled over all sample periods and for separate sample periods.

Table 4
Computing mobility with different consumption measures

	Number of transitions	Consumption per capita	Consumption per adult equivalent	Consumption filtered with demographic variables
	(1)	(2)	(3)	(4)
All years	8,107	0.4713 (0.0050)	0.5005 (0.0054)	0.5306 (0.0054)
1987-89	1,097	0.4702 (0.0146)	n.a.	n.a.
1989-91	1,914	0.4705 (0.0110)	0.5117 (0.0110)	0.5439 (0.0110)
1991-93	2,982	0.5029 (0.0089)	0.5273 (0.0089)	0.5594 (0.0089)
1993-95	3,211	0.4432 (0.0085)	0.4689 (0.0086)	0.4962 (0.0085)

Note: In column (1) the number of transitions for the row "All years" is 9,204. In 1987-89 the index cannot be computed because information on the number of children is missing in 1987. Standard errors are reported in parenthesis.

Table 5
The autocovariance matrix of consumption growth

	1989	1991	1993	1995
1989	0.1405 (0.0071)			
1991	-0.0443 (0.0081)	0.1398 (0.0056)		
1993	-0.0061 (0.0143)	-0.0643 (0.0070)	0.1748 (0.0064)	
1995	0.0121 (0.0162)	0.0049 (0.0055)	-0.0637 (0.0050)	0.1304 (0.0047)

Note: Standard errors are reported in parenthesis.

Table 6
Correcting for measurement errors

Year	Actual results		Montecarlo results		
			$\alpha = 0.05$	$\alpha = 0.1$	$\alpha = 0.12$
	$S(\widehat{\mathbf{P}})$	$s.e.(S(\widehat{\mathbf{P}}))$	$S(\widetilde{\mathbf{P}})$	$S(\widetilde{\mathbf{P}})$	$S(\widetilde{\mathbf{P}})$
1987-89	0.4702	0.0146	0.2572	0.3495	0.3749
1989-91	0.4705	0.0110	0.2650	0.3609	0.3882
1991-93	0.5029	0.0089	0.2616	0.3550	0.3825
1993-95	0.4432	0.0085	0.2564	0.3482	0.3772

Note: The Montecarlo simulation is described in Section 5.

Table 7
Computing mobility by demographic groups

Group		Group mobility		Diff. of means		
		$S(\widehat{\mathbf{P}})$	$s.e.(S(\widehat{\mathbf{P}}))$	$Gr.2$	$Gr.3$	$Gr.4$
1	North	0.5066	0.0079	2.51	0.92	
2	South	0.4782	0.0080		1.10	
3	Center	0.4937	0.0115			
1	Public	0.4470	0.0113	2.90		
2	Private	0.4865	0.0076			
1	Self-employed	0.5016	0.0183	1.75		
2	Employee	0.4681	0.0056			
1	Compulsory ed.	0.4881	0.0062	2.81	0.66	
2	High school	0.4524	0.0111		1.78	
3	College	0.4750	0.0188			
1	Born $\leq$ 1927	0.5002	0.0102	0.81	2.14	3.01
2	Born 1928-37	0.4881	0.0110		1.26	2.03
3	Born 1938-47	0.4690	0.0104			0.73
4	Born >1947	0.4589	0.0092			
1	No change in earners	0.4639	0.0055	4.62	4.80	
2	Pos. change in earners	0.5615	0.0204		0.51	
3	Neg. change in earners	0.5480	0.0166			

Note: The entries in the columns labelled $Gr.2$, $Gr.3$, and $Gr.4$ report the Z-statistic associated with the test that mobility for the group in that row equals mobility for groups 2, 3, and 4, respectively.

Figure 1: Mobility probabilities

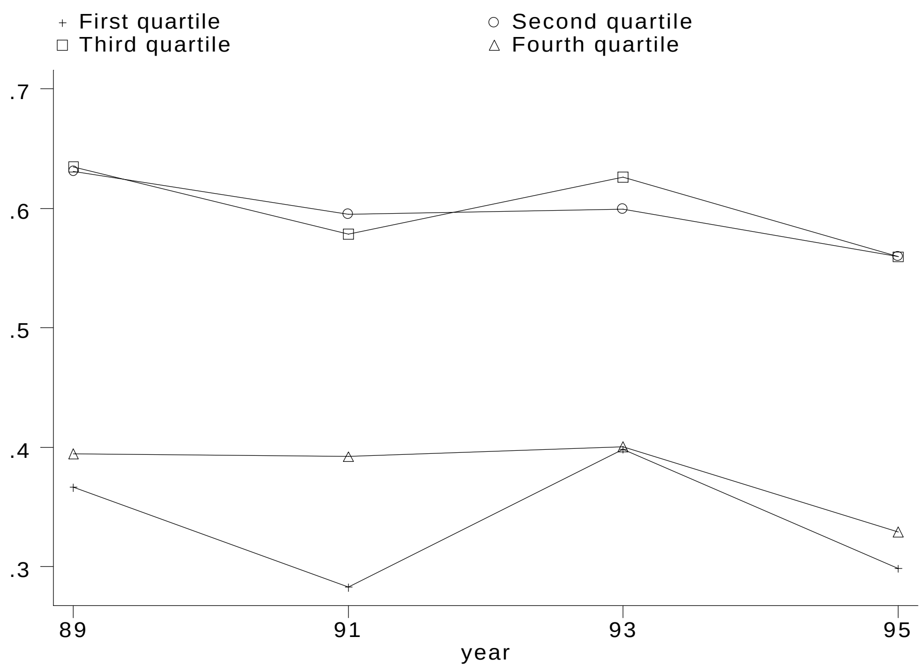

Figure 2: The distribution of annual working hours, 1989-1995